PROJECT MUSE®

Weakly Supervised Learning Creates a Fusion of Modeling Cultures

Chengliang Tang, Gan Yuan, Tian Zheng

Observational Studies, Volume 7, Issue 1, 2021, pp. 203-211 (Article)

Published by University of Pennsylvania Press

DOI: <https://doi.org/10.1353/obs.2021.0011>

➔ *For additional information about this article*

<https://muse.jhu.edu/article/799736>

Weakly Supervised Learning Creates a Fusion of Modeling Cultures

Chengliang Tang

ct2747@columbia.edu

Department of Statistics

Columbia University

1255 Amsterdam Avenue, New York, NY 10027

Gan Yuan

gy2277@columbia.edu

Department of Statistics

Columbia University

1255 Amsterdam Avenue, New York, NY 10027

Tian Zheng

tian.zheng@columbia.edu

Department of Statistics

Columbia University

1255 Amsterdam Avenue, New York, NY 10027

Abstract

The past two decades have witnessed the great success of the *algorithmic modeling* framework advocated by Breiman et al. (2001). Nevertheless, the excellent prediction performance of these black-box models rely heavily on the availability of *strong supervision*, i.e. a large set of accurate and exact ground-truth labels. In practice, strong supervision can be unavailable or expensive, which calls for modeling techniques under *weak supervision*. In this comment, we summarize the key concepts in weakly supervised learning and discuss some recent developments in the field. Using algorithmic modeling alone under a weak supervision might lead to unstable and misleading results. A promising direction would be integrating the *data modeling* culture into such a framework.

Keywords: Weakly Supervised Learning, Algorithmic Modeling, Data Modeling

As an important think piece to both the statistics and machine learning communities, Breiman et al. (2001) laid out the contrast of the two cultures in modeling thinking: *data modeling* and *algorithmic modeling*. It pointed out the limitations of data modeling and the opportunities and potentials of algorithmic modeling. Over the past two decades, Breiman et al. (2001)'s vision for algorithmic modeling has been validated by the rapid development and application of complicated yet effective algorithmic models, e.g., deep learning. Meanwhile, new challenges and opportunities are emerging everyday as we continue to deal with data with increasing size and complexity. In this comment, we offer a brief discussion of recent developments in the field of *weakly supervised learning* and discuss how it creates a need for *data modeling* thinking in an *algorithmic modeling* framework.

Following the taxonomy introduced in Breiman et al. (2001), the *data modeling* culture refers to methods that explicitly assume a stochastic model for data generation. Often,

methods of this culture have shallow structures and are easy to interpret. Typical examples include linear regression, logistic regression, to name a few. The validity of such methods is backed by the probabilistic properties of their outputs, such as goodness-of-fit tests and residual analyses. In contrast, the *algorithmic modeling* culture aims to learn the complex and unknown nature of true data generation mechanisms through “black-box” algorithms. Typical examples of this culture include decision trees, support vector machines (SVM), and neural networks (NN). The training and evaluation of these algorithms are guided by predictive accuracy.

The past two decades have witnessed the rapid expansion and success of the algorithmic modeling culture. From self-driving cars (Bojarski et al., 2016) to virtual assistants (Devlin et al., 2018), complicated algorithmic models such as *deep neural networks* (DNNs) have demonstrated their potential for leveraging today’s big data and affordable high-performance computational resources in producing predictions that are comparable to human performance. However, training such algorithms to attain impressive performance relies heavily on a large volume of training data with high-quality labels (see Figure 1), which are often expensive or even unavailable in many real-world applications. In particular, such a strong supervision becomes substantially scarcer in application domains that are more specialized, such as healthcare (Miotto et al., 2018) and ecological studies (Christin et al., 2019; Tang et al., 2021), where domain expertise is vital in data labeling. As a result, practical challenges due to the lack of strong supervision in many real-world applications significantly limit the applicability and generalization of algorithmic models.

Weakly supervised learning (WSL) (Zhou, 2018) addresses the more realistic setting when supervision is available but weak under various practical scenarios. It expands the reach of conventional supervised learning and has garnered a lot of interests in applications (e.g. Jorgensen et al., 2008; Oquab et al., 2015; Peyre et al., 2017). In algorithmic modeling, strong supervision comes from a large set of accurately labelled data. Such a supervision may be weakened in approximately three ways: incomplete supervision, inexact supervision, and inaccurate supervision (Zhou, 2018).

Figure 1: Strong supervision drives performance improvement in supervised learning.

Let $\mathbf{X}$ be the input features. Let $\mathbf{Y}$ be the outcome of interest. When $\mathbf{Y}$ ’s values are available in the training data as labels, they provide strong supervision for the algorithmic modeling, through a loss function $L(\mathbf{Y}, \hat{\mathbf{Y}})$ (Figure 1). In practice, the exact and accurate values of $\mathbf{Y}$ are often unavailable in the training data. Instead, let $\tilde{\mathbf{Y}}$ be the observed (weak) labels in the training data. Here, we introduce a unified notation, $\mathbf{W}$, for the generating

mechanism of the weakened supervision $\tilde{\mathbf{Y}}$. Using the above notation, the framework of WSL is summarized in Figure 2. WSL shares the same learning goal with methods of the algorithmic modeling culture in Breiman et al. (2001), that is, train a function $f(\mathbf{X})$ such that $\mathbf{Y}$ can be accurately predicted or approximated by $f(\mathbf{X})$. Here, the challenges arise mostly from the lack of strong labels $\mathbf{Y}$ and the need to create effective supervision based on the observed $\tilde{\mathbf{Y}}$.

Figure 2: Weakly supervised learning poses new challenges for the algorithmic modeling framework.

Directly applying algorithmic modeling to data with weakened training labels *without* considering the weak supervision generating mechanism $\mathbf{W}$ could lead to results that are unstable and overfitted (Frénay and Verleysen, 2013; Van Engelen and Hoos, 2020). Take semi-supervised learning as an example, which can be thought of as a special case of *incomplete supervision* (Zhou, 2018). Most algorithms in the field of semi-supervised learning rely on the assumption that the labels are missing complete at random. When this assumption is violated in real data, semi-supervised learning algorithms may actually degrade the learning performance, compared to applying supervised learning methods directly on the labeled portion of the dataset (Zhu, 2008). Another example is training DNNs with noisy training labels, i.e., *inaccurate supervision*. Zhang et al. (2017) provided empirical results showing that DNNs can fit training data with randomly shuffled labels arbitrarily well. Not surprisingly, the generalization performance of the trained DNNs on test sets was no better than random guessing. Even for the relatively shallow tree ensemble models, numerical experiments have shown that the adaptive boosting algorithm (AdaBoost) would disproportionately focus on learning mislabeled instances when label noises exist (Dietterich, 2000). Therefore, an algorithmic modeling framework under weak supervision needs to explicitly acknowledge the weakening mechanism $\mathbf{W}$.

The entire promise of WSL lies within the assumption that the weak labels $\tilde{\mathbf{Y}}$ in the training data carry partial information of $\mathbf{Y}$ through the weak supervision generating mechanism $\mathbf{W}$. Most current methods in WSL assume that the mapping from $\mathbf{Y}$ to $\tilde{\mathbf{Y}}$ by $\mathbf{W}$ is independent of the features $\mathbf{X}$ and the true labels $\mathbf{Y}$. A more realistic scenario, however, would be that the mechanism of $\mathbf{W}$ may be dependent of both $\mathbf{X}$ and $\mathbf{Y}$. Consider the joint distribution of the observed weakened labels $\tilde{\mathbf{Y}}$ and the features $\mathbf{X}$,

$$P(\tilde{\mathbf{Y}}, \mathbf{X}) = \sum_{\mathbf{Y}} P(\tilde{\mathbf{Y}}|\mathbf{Y}, \mathbf{X})P(\mathbf{Y}|\mathbf{X})P(\mathbf{X}). \quad (1)$$

As shown in Figure 2, our learning goal remains to be fitting $P(\mathbf{Y}|\mathbf{X})$, the unknown *data generating mechanism*, with a model $f(\mathbf{X})$, even when we lack direct observations of $\mathbf{Y}$. To allow information in $\tilde{\mathbf{Y}}$ be passed onto the learning of $f(\mathbf{X})$, it is critical to model $\mathbf{W}$ in Figure 2, i.e., $P(\tilde{\mathbf{Y}}|\mathbf{Y}, \mathbf{X})$ on the right-hand side of Equation (1). In Figure 3, we introduce an overly generalized notation $g(\tilde{\mathbf{Y}}|\mathbf{Y}, \mathbf{X})$ to encapsulate models and approaches for the mechanism $\mathbf{W}$. In practice, characterizing $P(\tilde{\mathbf{Y}}|\mathbf{Y}, \mathbf{X})$ could be challenging as information is often scarce. In fact, without additional information beyond the training data, it is not possible to effectively leverage the weak supervision that is offered by $\tilde{\mathbf{Y}}$ (Frénay and Verleysen, 2013; Zhou, 2018). In the weakly supervised learning literature, additional information for constructing $g(\tilde{\mathbf{Y}}|\mathbf{Y}, \mathbf{X})$ has been introduced in the form of assumptions on $\mathbf{W}$ and/or small sets of data with observed $\mathbf{Y}$. This is primarily motivated by the need for transparency and interpretability for $g(\tilde{\mathbf{Y}}|\mathbf{Y}, \mathbf{X})$ to incorporate prior knowledge into “end-to-end” modeling frameworks. In other words, it is desirable to have the modeling of $\mathbf{W}$ be “assumption-driven” rather than “data-driven” or “accuracy-driven”, which creates a role for the *data modeling* culture within an *algorithmic modeling* framework (Figure 3). In the weakly supervised learning literature, there has been some progress made to address each type of weak supervision.

Figure 3: Weakly supervised learning requires a fusion of data modeling thinking into the algorithmic modeling framework.

For *incomplete supervision* where labels are only available for a small subset of training data, active learning algorithms (Settles, 2009) attempt to better extract label information by “actively” asking an “oracle” (e.g., a human annotator) for queries of selected unlabeled instances. This framework has been widely used in image classification (Joshi et al., 2009; Kapoor et al., 2007; Li and Guo, 2013). Assuming the existence of an “oracle”, the key component of active learning is to choose the most “valuable” instance to query. To this end, measures of informativeness and representativeness of individual observations have been proposed (Settles, 2009). For example, Bayesian active learning methods estimate the expected improvement of each instance query through nonparametric models such as Gaussian process and Monte Carlo estimations (e.g. Gal et al., 2017; Kapoor et al., 2007;

Roy and McCallum, 2001). As another approach to incomplete supervision, semi-supervised learning algorithms (Chapelle et al., 2009; Zhu, 2005) utilize the unlabeled training data as well as labeled data to improve prediction accuracy. Transductive methods were proposed to obtain label prediction for unlabelled data points (Van Engelen and Hoos, 2020), including the use of probabilistic models, such as Markov random fields and Gaussian random fields, for label assignments (e.g. Shental and Domany, 2005; Wu et al., 2012; Zhu and Ghahramani, 2002).

Inexact supervision addresses the situation where the given labels are at coarser scales than desired. For example, in many real-world object segmentation tasks, only image-level training labels are available, while the task is to localize each object. Multi-instance learning (Zhou and Zhang, 2007) was such an example with a *bag-of-instances* setup: instances x_{ij} are organized in bags X_j , and the labels in the training set are only given at the bag level. A common assumption for this task is that the bag-level class probability is the maximum of all the instance-level class probabilities within the bag. This assumption bridges the gap between instance predictions and observed bag labels. Another example is the *concept labeling* method (Chenthamarakshan et al., 2011), which assumes a soft bag-instance structure. In their Bayesian modeling framework, each document (instance) $\mathbf{X}$ has a distribution $P(\mathbf{V}|\mathbf{X})$ over the concepts of the ontology $\mathbf{V}$ (bag). It is assumed the outcome variable of interest $\mathbf{Y}$, categories, is conditionally independent of document content $\mathbf{X}$, when conditioning on the ontology concept $\mathbf{V}$. Consider $P(\mathbf{Y}|\mathbf{X}) = \sum_{\mathbf{V}} P(\mathbf{Y}, \mathbf{V}|\mathbf{X}) = \sum_{\mathbf{V}} P(\mathbf{Y}|\mathbf{V})P(\mathbf{V}|\mathbf{X})$. As a result, by separately modeling the document-to-concept distribution $P(\mathbf{V}|\mathbf{X})$ and the concept-to-class distribution $P(\mathbf{Y}|\mathbf{V})$, the instance-level document label predictions $P(\mathbf{Y}|\mathbf{X})$ can be obtained.

Inaccurate supervision concerns the situation where labels are a noisy version of the ground truth. To learn with noisy labels, many algorithms make the assumption that the noises are randomly generated. Brodley and Friedl (1999) proposes to first identify the potentially mislabeled instances and perform label correction. Northcutt et al. (2021) proposed the *Confidence Learning* framework that iteratively determines which labels are more likely to be the contaminated ones, based on an estimated joint distribution of true label $\mathbf{Y}$ and observed label $\tilde{\mathbf{Y}}$. The *data programming* approach proposed by Ratner et al. (2016) is a paradigm for integrating noisy labels from multiple sources, and deriving a better training set using a dependency graph that incorporates different assumptions on the weak supervision generating mechanisms.

For any WSL framework, optimizing the generalization performance of the learned model $f(\mathbf{X})$, for $P(\mathbf{Y}|\mathbf{X})$, remains the main goal. However, it is important to consider the practical issues caused by the imperfection of available data and construct “end-to-end” learning frameworks that take raw training data and deliver reliable final models. In this comment, we argue that the scarcity of strong supervision in many real-world applications calls for a fusion of modeling cultures that allow creative combinations of assumption-driven and data-driven approaches. There are many open problems and challenges that remain to be further explored. In particular, in many real-world learning tasks, all the above weak supervision scenarios may apply at the same time (e.g., noisy and inexact labels are only available on a small subset, as seen in Tang et al. (2021)). Most of the existing weakly supervised learning methods focus only on a single type of weak supervision. As a result, the weak supervision generation mechanism $P(\tilde{\mathbf{Y}}|\mathbf{Y}, \mathbf{X})$ is usually over-simplified in practice.

Much of the statistical literature from the data modeling culture, e.g., robust statistics and methods for missing data, may find application in end-to-end workflows of weakly supervised learning. In addition, many current methods in WSL incorporate assumptions on the weak supervision in an ad hoc fashion. For the same reasons that have led to the lack of strong supervision in the training data, it is also impractical to assume that one can validate the learning framework using prediction accuracy on test data alone. Systematic model checking with respect to the weak supervision generating mechanism $\mathbf{W}$ is needed.

References

- Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Praseoon Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- Leo Breiman et al. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical science*, 16(3):199–231, 2001.
- Carla E Brodley and Mark A Friedl. Identifying mislabeled training data. *Journal of artificial intelligence research*, 11:131–167, 1999.
- Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews]. *IEEE Transactions on Neural Networks*, 20(3):542–542, 2009.
- Vijil Chenthamarakshan, Prem Melville, Vikas Sindhwani, and Richard D Lawrence. Concept labeling: Building text classifiers with minimal supervision. In *Twenty-Second International Joint Conference on Artificial Intelligence*, 2011.
- Sylvain Christin, Éric Hervet, and Nicolas Lecomte. Applications for deep learning in ecology. *Methods in Ecology and Evolution*, 10(10):1632–1644, 2019.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Thomas G Dietterich. An experimental comparison of three methods for constructing ensembles of decision trees: Bagging, boosting, and randomization. *Machine learning*, 40(2):139–157, 2000.
- Benoît Frénay and Michel Verleysen. Classification in the presence of label noise: a survey. *IEEE transactions on neural networks and learning systems*, 25(5):845–869, 2013.
- Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *International Conference on Machine Learning*, pages 1183–1192. PMLR, 2017.
- Zach Jorgensen, Yan Zhou, and Meador Inge. A multiple instance learning strategy for combating good word attacks on spam filters. *Journal of Machine Learning Research*, 9(6), 2008.
- Ajay J Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos. Multi-class active learning for image classification. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2372–2379. IEEE, 2009.
- Ashish Kapoor, Kristen Grauman, Raquel Urtasun, and Trevor Darrell. Active learning with Gaussian processes for object categorization. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007.

- Xin Li and Yuhong Guo. Adaptive active learning for image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 859–866, 2013.
- Riccardo Miotto, Fei Wang, Shuang Wang, Xiaoqian Jiang, and Joel T Dudley. Deep learning for healthcare: review, opportunities and challenges. *Briefings in bioinformatics*, 19(6):1236–1246, 2018.
- C. G. Northcutt, Lu Jiang, and Chuang I. L. Confident learning: Estimating uncertainty in dataset labels. *arXiv preprint arXiv: 1911.00068v4*, 2021.
- Maxime Oquab, Léon Bottou, Ivan Laptev, and Josef Sivic. Is object localization for free?-weakly-supervised learning with convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 685–694, 2015.
- Julia Peyre, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Weakly-supervised learning of visual relations. In *Proceedings of the IEEE international conference on computer vision*, pages 5179–5188, 2017.
- Alexander J Ratner, Christopher M De Sa, Sen Wu, Daniel Selsam, and Christopher Ré. Data programming: Creating large training sets, quickly. In *Advances in neural information processing systems*, pages 3567–3575, 2016.
- Nicholas Roy and Andrew McCallum. Toward optimal active learning through monte carlo estimation of error reduction. *ICML, Williamstown*, pages 441–448, 2001.
- Burr Settles. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009.
- Noam Shental and Eytan Domany. Semi-supervised learning—a statistical physics approach. In *“Learning with Partially Classified Training Data”—ICML05 workshop*. Citeseer, 2005.
- Chengliang Tang, María Uriarte, Helen Jin, Douglas Morton, and Tian Zheng. Large-scale, image-based tree species mapping in a tropical forest using artificial perceptual learning. *Methods in Ecology and Evolution*, 2021.
- Jesper E Van Engelen and Holger H Hoos. A survey on semi-supervised learning. *Machine Learning*, 109(2):373–440, 2020.
- Xiao-Ming Wu, Zhenguo Li, Anthony Man-Cho So, John Wright, and Shih-Fu Chang. Learning with partially absorbing random walks. In *NIPS*, volume 25, pages 3077–3085, 2012.
- C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning requires rethinking generalization. *International Conference on Learning Representations (ICLR)*, 2017.
- Zhi-Hua Zhou. A brief introduction to weakly supervised learning. *National science review*, 5(1):44–53, 2018.

- Zhi-Hua Zhou and Min-Ling Zhang. Multi-instance multi-label learning with application to scene classification. In *Advances in neural information processing systems*, pages 1609–1616, 2007.
- X. Zhu. Semi-supervised learning literature survey. *Technical Report. 1530, University of Wisconsin Madison*, 2008.
- Xiaojin Zhu and Zoubin Ghahramani. Towards semi-supervised classification with markov random fields. 2002.
- Xiaojin Jerry Zhu. Semi-supervised learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2005.